Oversight Committee Members

(Anna Hasenfratz) anna@eotvos.colorado.edu
(Greg Kilcup) kilcup@pacific.mps.ohio-state.edu
(Julius Kuti) jkuti@ucsd.edu
(Rob Pennington) robp@ncsa.uiuc.edu
(Ralph Roskies) roskies@psc.edu
(Terry Schalk) tas@scipp.ucsc.edu
(Steven Gottlieb) sg@fuji.physics.indiana.edu

(Robert Sugar) sugar@physics.ucsb.edu *

* (Bob is on the email list, but says he is not an ex officio committee member.)

To email all the Oversight Committee members, send email to:

oversight at fuji.physics.indiana.edu

Oversight Committee Charge

The Oversight Committee will report to the Principal Investigators. It will:

1) Review progress in implementing the plans of the collaboration.

2) Review plans for the development and acquisition of software and hardware.

3) Make recommendations regarding alternative approaches or new directions for the project.

The Oversight Committee will meet in-person at least twice a year. At its meetings it will be briefed by the Software Coordinator and the software and hardware developers regarding their plans and progress, and then make recommendations regarding them to the Principal Investigators.

Between meetings the Chair of the Oversight Committee will maintain regular contact with the software and hardware developers and the Software Coordinator, and will keep his Committee informed of developments. The Oversight Committee will hold conference calls between its meetings as deemed appropriate by its Chair. It will provide reports to the Principal Investigators on a quarterly basis.

The schedule of contacts and meetings will be arranged so that the Oversight Committee can review on-going progress and planning and provide timely advice before important implementation or procurement decisions are taken.

Letter from Oversight Committee to Executive Committee

On October 21, the Oversight Committee of the Lattice Gauge Theory SciDAC project met to consider the latest proposal.

The Oversight Committee is deeply concerned that the current proposal has dropped from the timeline an evaluation of the 128 cpu machine before construction of the 1.5 Teraflop machine is substantially underway. In February, the collaboration members at the All-Hands Meeting expressed strong support for the statement:

"We believe that in order to have a credible proposal for a 1 TF QCDOC prototype it will be necessary for the proposal to detail a two phase project with appropriate milestones achieved in the first phase. In phase one, on a 100–200 GF platform, the hardware operation and the efficiency of the high level QCD software which is crucial for the physics of the user community must be verified. After this milestone, the 1 TF QCDOC prototype would be built to verify favorable scaling properties and to do substantial physics research. This approach provides a reasonable safeguard against unforseen serious hardware and software difficulties.

"In summary, we support the preparation of a modified proposal with the understanding that both prototypes will be available for code development and production running by the broad community of SciDAC grant collaborators as quickly as possible. We believe that bypassing a smaller prototype platform for hardware and software tests would involve unwanted higher risk. After successful verification, with secured funding in hand, the 1 TF platform should be constructed with agressive scheduling."

At some point it needs to be demonstrated that QCDOC works and how useful it is for the general community. The sooner that this is done, the sooner the collaboration can make a decision about the future mix of hardware that we will purchase. At each stage of development it is crucial to demonstrate the usefulness of the hardware. The same issue of delay in schedule vs. verification of hardware is likely to come up regarding a future proposal. It is important to make sure that we do not establish a pattern of forging ahead with requests without demonstrating that we are reaching our goals in terms of performance and usefulness.

The oversight committee wishes to make it clear that we feel the neglect of appropriate safeguards and milestones is unwise and is not something that can be repeated in future proposals.

The oversight committee respectfully requests a statement from the executive committee that addresses the following two questions: 1) Why did the executive committee feel that it was necessary to deviate from the plan accepted by the collaboration at the All-Hands meeting at Jefferson Laboratory? 2) How can the executive committee assure that future proposals require achievement of milestones prior to expansion of computing facilities?

The committee also expressed concern about the length of time it is taking to open to the collaboration the cluster hardware that was purchased this summer. It is important to open these facilities as soon as possible.

Response from Executive Committee

Dear Members of the Oversight Committee,

Thank you very much for your input to the 2003 Hardware Proposal. As you requested, we will address the two questions raised in your report:

1) Why did the executive committee feel that it was necessary to deviate from the plan accepted by the collaboration at the All-Hands meeting at Jefferson Laboratory?
2) How can the executive committee assure that future proposals require achievement of milestones prior to expansion of computing facilities?

First, we want to emphasize that we are in complete agreement with the recommendation of the Oversight Committee and the full collaboration that appropriate milestones be formulated and achieved for each phase of the project before moving on to the next phase. In the case of the QCDOC we have established milestones to be achieved by the hardware, and are in the process of developing ones for the software. The former are set out in detail in the proposal, and a very preliminary draft of software milestones can found at the URL physics.bu.edu/ brower/tests/. In our current plan, the hardware milestones must be achieved on the ASIC chips and on the 128 node prototype before construction is begun on the 1.5 Tflops development machine. In addition, sufficient software milestones must be met to demonstrate that important physics can be done on the development machine. The achievement of these milestones will determine whether or not the development machine is built. Tests of the SciDAC software and the software of individual groups will begin on the prototype and will be

completed on the development machine. The decision as to whether or not to build a large QCDOC at BNL will be based on the performance of this software on the prototype and the development machines, and on an assessment as to whether the QCDOC is a suitable platform for serving the broad U.S. lattice gauge theory community. Thus, in this approach the prototype will be used primarily to test the hardware, and the development machine to test the community software. What has changed since the JLab meeting is not our commitment to milestones, but a re-assignment of the platforms on which some of the software tests will be run. These changes came about in part because the Columbia/IBM team was not able to keep up with the very aggressive schedule for the QCDOC presented at the JLab meeting, and in part because the SciDAC software will not be completed in time to be fully tested on the prototype without seriously delaying the hardware effort. We believe that we have arrived at a plan which properly balances the risks of moving forward too rapidly with those of losing momentum by acting too slowly. However, it is now clear that we did not keep you adequately informed regarding the evolution in our thinking, and that we did not explain the testing process clearly enough in the draft proposal. We apologize for the former, and have made changes to the proposal in an attempt to correct the latter. (Please see pages 12 and 44-46 of the draft currently on the web).

It may be helpful to begin by briefly outlining the tests we currently envision, and the purpose of each of them. We propose to perform tests on three QCDOC hardware platforms:

1) Tests of the ASIC: The ASIC will be tested by running the full Columbia Physics System Code, as well as the optimized improved staggered inverter. The purpose of these tests is to verify that the ASIC works as designed. If it does not, then the problems will be addressed and the chip re-spun. If, as we expect, the ASIC does work as designed, the Columbia group will proceed to build the 128 node prototype.
2) Tests of the 128 Node Prototype: The tests of the prototype will address four issues: a) The performance and stability of the QCDOC hardware b) The performance and stability of the QCDOC operating system c) Scaling of the Columbia/UKQCD code to 128 nodes d) Performance of early SciDAC software
3) Tests of the 1.5 Tflops Development Machine: These tests will address in detail

a) The performance of the SciDAC code and the code of individual groups, such as MILC, that have been brought into confority with the SciDAC standards.
b) The suitability of the QCDOC to be a computing platform for the entire community.

Now to address your specific questions:

1) As indicated above, the major tests of the QCDOC hardware will be carried out on the ASIC chips. These must be passed before construction of the 128 node prototype begins. We do not believe that a significant risk is entailed by ordering ASIC chips for the prototype or the development machines once the ASIC have been successfully tested, because the design and construction of the mother and daughter boards is relatively straightforward, and the Columbia group has considerable experience in this area. What has changed since the JLab meeting is that much of the detailed testing of the SciDAC code, and the code of individual groups, has been moved from the 128 node prototype to the 1.5 Tflops development machine. The alternative was a considerable delay in the construction of the development machine (and the subsequent BNL machine), thereby risking a loss of momentum for the project as a whole. We believe that the risk taken in beginning construction of the development machine before fully testing software on the prototype is modest and worthwhile for the following reasons. First, it is straightforward to accurately predict the performance for QCD of a parallel machine with a mesh architecture once the performance of a single node is known. Thus, based on tests run on the simulator, we are confident that if the ASIC does work as designed, then the QCDOC will yield the very strong performance for the Columbia/UKQCD code indicated by Peter Boyle in his Lattice 2002 presentation (arxiv.org/abs/hep-lat/0210034). Furthermore, we have set up a schedule which we are confident will produce code that will generate Asqtad lattices with high efficiency by the time the development machine goes into operation. (Recall that the Asqtad action is the only one designated by the Scientific Program Committee for early use on the QCDOC that is not available in the Columbia Physics System code). Thus, at the very minimum, the 1.5 Tflops development machine will greatly enhance the capability of the U.S. lattice gauge theory community to produce realistic lattices for the study of QCD.

We submit that these lattices will be valuable enough by themselves to justify the construction of the 1.5 Tflops machine. Furthermore, we believe that it is essential to perform the tests of the development machine set out above before building a large QCDOC at BNL. To determine whether that QCDOC will truly meet the needs of the U.S. community will require experimentation on a computer with capabilities approximating those we propose for the development machine. In our opinion, jumping from the 128 node prototype to a 10 Tflops machine without demonstrating that the QCDOC will meet the needs of the community would be much riskier than the course we propose. The fact that the UKQCD and Riken groups have come to the same conclusion adds weight to this assertion. Finally, although we do not in any way want to minimize the expenditure of $1.5M, we note that this amount is less than has already been committed to the construction of prototype clusters through our SciDAC grant.

2) Again, we agree that it is very important to have well defined milestones which must be achieved before large new facilities are built. We are in the process of developing detailed milestones for the QCDOC and clusters, and will seek your advice regarding them. With respect to the QCDOC in particular, we are committed to successfully completing tests on the 1.5 Tflops development machine before beginning construction of a multi-teraflops machine at BNL.

We hope that we have answered your questions. We would appreciate any further advice that you have regarding the proposal.

Sincerely,

The Executive Committee